

SANGMIN SEO

Mathematics and Computer Science (MCS) Division, Argonne National Laboratory (ANL)

Address: 9700 South Cass Avenue, Bldg. 240, Rm. 3155, Lemont, IL 60439, USA

Phone: +1-630-252-1363 **Email:** sseo@anl.gov **Web:** <http://www.mcs.anl.gov/~sseo/>

1. Resume

1.1. Education

Ph.D. in Electrical Engineering and Computer Science SEP 2007 - AUG 2013

- Seoul National University, Seoul, Korea
- Thesis: *Enhancing Performance Portability of OpenCL for Multicore CPUs*
- Advisor: Prof. Jaejin Lee
- GPA: 4.12/4.3

B.S. in Computer Science and Engineering MAR 2001 - AUG 2007

- Seoul National University, Seoul, Korea
- GPA: 3.77/4.3 (Cum Laude)

1.2. Appointments

Argonne National Laboratory, Lemont, IL, USA

- Assistant Computer Scientist APR 2016 - present
- Postdoctoral Researcher APR 2014 - APR 2016

ManyCoreSoft, Seoul, Korea

- Chief Executive Officer (CEO) JUN 2012 - FEB 2014

1.3. Research Interests

High-performance computing, parallel programming models, compilers, and parallel runtime systems

1.4. Honors and Awards

Bronze Prize FEB 2009

- In the 15th SAMSUNG HumanTech Thesis Competition
- Paper: Software-Managed Caches for Multicores with Local Memory

Doctoral Fellowship SEP 2008 - AUG 2012

- Korea Foundation for Advanced Studies

Graduate Student Fellowship SEP 2008 - FEB 2010

- Sinyang Cultural Foundation, Korea

ChoDeokHee Scholarship MAR 2007

- Dongwon Enterprise, Korea

First Prize APR 2005

- In Office Designer Category of Microsoft Imagine Cup 2005 Local Competition
- With Subeom Lee and Kyoungjun Lee
- Designed and implemented a management solution for art museums, called *Curatee*, on the Microsoft Office, web service, and Windows Mobile platform

Undergraduate Student Scholarship SEP 2001 - MAR 2004, MAR 2007

- Seoul National University, Seoul, Korea

1.5. Membership in Professional Societies

- [1] Professional Member, IEEE Computer Society (2013 - present)
- [2] Professional Member, IEEE (2013 - present)
- [3] Professional Member, ACM (2013 - present)
- [4] Student Member, IEEE Computer Society (2009 - 2013)
- [5] Student Member, IEEE (2009 - 2013)
- [6] Student Member, ACM (2009 - 2013)

2. Experience

2.1. Research Experience

Assistant Computer Scientist, *Argonne National Laboratory*

APR 2016 - *present*

Postdoctoral Researcher, *Argonne National Laboratory*

APR 2014 - APR 2016

- **ARGO**: An Exascale Operating System and Runtime Research Project
 - Working on the concurrency runtime, called **Argobots** - a lightweight low-level threading and tasking infrastructure, and leading the development of Argobots
 - Collaborating with several teams of parallel programming models, such as OpenMP, MPI, Charm++, Cilk, PaRSEC, TASCEL, OmpSs, XcalableMP, and Mercury, to integrate Argobots into their runtimes
- **MPICH**: A High-Performance, Portable Implementation of the MPI Standard
 - Participating in the development and maintenance of MPICH as a member of the MPICH team
 - Served as the release manager for MPICH 3.1.3
 - Implemented nonblocking collective (NBC) I/O operations, which are included in the latest MPI 3.1 standard
 - Conducted a research on the interoperation between the MPI runtime and user-level threading models; and developed threading layers inside MPICH to support Argobots and Qthreads
 - Studied thread synchronization issues in multi-threaded communication and devised work-driven thread synchronization techniques as well as locality-preserving locking mechanisms
- **BOLT**: OpenMP over Lightweight Threads
 - BOLT targets a high-performing OpenMP implementation, specialized for nested or fine-grain parallelism, and utilizes a lightweight threading model (currently Argobots) as its underlying threading mechanism.
 - Leading the project as PI and developing the runtime and compiler based on Intel's open-source OpenMP runtime and LLVM/Clang, respectively

Graduate Student, *Seoul National University*

SEP 2007 - AUG 2013

- As a member of **Multicore Computing Research Laboratory** and **Center for Manycore Programming**, I participated in several research projects:
- **SnuCL**: An OpenCL Framework for Heterogeneous CPU/GPU Clusters
 - An OpenCL framework that extends the original OpenCL semantics to the heterogeneous cluster environment. With SnuCL, OpenCL applications written for a single node can run on the heterogeneous cluster without any modifications.
 - Developed the OpenCL compiler in SnuCL
 - Also devised an automatic OpenCL work-group size selection technique for multicore CPUs
- **Chundoong**: A Low-Cost Energy-Efficient Heterogeneous Supercomputer
 - A self-made, 60-node (56 compute nodes + 4 storage nodes) heterogeneous CPU/GPU cluster equipped with a self-designed water cooling system
 - Ranked 277th in the TOP500 list and 32nd in the Green500 list of November 2012
 - Participated in building the Chundoong cluster and optimizing the kernel code of LINPACK benchmark
- **SNU NPB Suite**: A set of the NAS Parallel Benchmarks (NPB) implemented in C, OpenMP C, and OpenCL
 - Also includes an OpenCL version using multiple OpenCL compute devices
 - Was the main developer for this benchmark suite and ported most NPB applications
- **SNU-SAMSUNG OpenCL Framework**
 - An OpenCL framework that targets heterogeneous multicore architectures with local memory, such as Cell BE processors. It also supports ARM processors and DSPs.

- Consists of the runtime and OpenCL C source-to-source translator
- Solely developed the OpenCL C-to-C translator using Clang, which is a C/C++ front-end of LLVM
- **SFMalloc**: A lock-free and mostly synchronization-free dynamic memory allocator for manycores
 - SFMalloc never uses any synchronization for common cases and uses only lock-free synchronization mechanisms for uncommon cases.
 - Solely designed and implemented SFMalloc
- **Extended Set-index Cache (ESC)**
 - A software-managed cache that has the benefits of set-associative cache and fully associative cache. Its tag search speed is comparable to that of the set-associative cache and its miss rate is comparable to that of the fully associative cache. ESC was used in COMIC, COMIC++, and other research projects.
 - Designed and implemented ESC
- **COMIC**: A Coherent Shared Memory Interface for Cell BE
 - A software shared virtual memory (SVM) that supports SPMD-style parallel programming for heterogeneous multicores with small local memory such as Cell BE processors
 - Played a role of the main developer
- **COMIC++**: A Software SVM System for Heterogeneous Multicore Accelerator Clusters
 - Participated in the design of memory coherence and consistency protocol and modified COMIC as an internal node protocol
- **Software Transactional Memory (STM) for Multicore Accelerator Architectures**
 - An STM system that targets multicore accelerator architectures for compute-intensive applications
 - Helped the design of STM in the heterogeneous multicore architecture and implemented an efficient dynamic memory manager used in the STM system

2.2. Industry Experience

Chief Executive Officer (CEO)

JUN 2012 - FEB 2014

- [ManyCoreSoft](#), Seoul, Korea
- Built a heterogeneous CPU/GPU supercomputer named “Chundoong”
 - A self-made, 60-node (56 compute + 4 storage) cluster equipped with a self-designed water cooling system
 - In charge of core development together with Seoul National University
 - Ranked 277th in the TOP500 list and 32nd in the Green500 super computer list of Nov 2012
- Developed a GPU accelerated computational finance system
 - Core part of the commercial product of Koscom, a leading Korean IT company in finance sector
 - Led the project, mainly developed the parallel processing algorithm using OpenCL and OpenMP
 - The average performance speedup was about 4x, up to 16x
- Obtained two research grants from the Korean government
 - *Development of the Big Data Processing Platform based on Manycore Performance Acceleration*
 - *Chundoong-S: High-performance Accelerator Supercomputing Solution*
 - Led the overall project management, e.g., writing proposals, presentation, winning research fund
- Built two CPU/GPU clusters for National Institute for Mathematical Sciences in Korea

Military Service

MAR 2004 - JUL 2006

- [Korea Optron Corporation](#), Gwangju, Korea
- Finished military duty at a company chosen by the Military Manpower Administration as an alternative to serving in the military
- Developed a wiki engine using PHP and MySQL
- Participated in the development of optical power meter that measures the energy in an optical signal and optical transceiver module that transmits and receives data using optical fiber

Winter Intern

JAN 2004

- [Samsung Electronics](#), Suwon, Korea
- Experience with mobile phone development

Summer Intern

JUL 2003

- [Samsung Electronics](#), Suwon, Korea

- Experience with mobile phone development

Author

SEP 2001 - DEC 2001

- ETOOS Publishing Company, Seoul, Korea
- Wrote chapters about differential and integral calculus of high school mathematics reference book

3. Research Products

3.1. Refereed Journal Articles

- [1] Hanqi Guo, Wenbin He, Sangmin Seo, Han-Wei Shen, and Tom Peterka. *Extreme-Scale Stochastic Particle Tracing for Uncertain Unsteady Flow Analysis*. IEEE Transactions on Parallel and Distributed Systems (TPDS), 2016 (*under review*).
- [2] Thanh Tuan Dao, Jungwon Kim, Sangmin Seo, Bernhard Egger, and Jaejin Lee. *A Performance Model for GPUs with Caches*. IEEE Transactions on Parallel and Distributed Systems (TPDS), Vol. 26, No. 7, pp. 1800–1813, Jul. 2015, DOI:10.1109/TPDS.2014.2333526.

3.2. Publications in Refereed Conference and Workshop Proceedings

- [1] Philip Carns, John Jenkins, Sangmin Seo, Shane Snyder, Robert B. Ross, Charles D. Cranor, Scott Atchley, Torsten Hoefler. *Enabling NVM for Data-Intensive Scientific Services*. In INFLOW '16: Proceedings of the 4th Workshop on Interactions of NVM/Flash with Operating Systems and Workloads, Savannah, GA, USA, Nov. 2016.
- [2] Adrián Castelló, Antonio J. Peña, Sangmin Seo, Rafael Mayo, Pavan Balaji, and Enrique S. Quintana-Ortí. *A Review of Lightweight Thread Approaches for High Performance Computing*. In Cluster '16: Proceedings of the 2016 IEEE International Conference on Cluster Computing, Taipei, Taiwan, Sep. 2016. (39/162, 24.1%)
- [3] Jintao Meng, Sangmin Seo, Pavan Balaji, Yanjie Wei, Bingqiang Wang, and Shenzhong Feng. *SWAP-Assembler 2: Optimization of De Novo Genome Assembler at Extreme Scale*. In ICPP '16: Proceedings of the 45th International Conference on Parallel Processing, Philadelphia, PA, USA, Aug. 2016, DOI:10.1109/ICPP.2016.29.
- [4] Daniel Ellsworth, Tapasya Patki, Swann Perarnau, Sangmin Seo, Abdelhalim Amer, Judicael Zounmevo, Rinku Gupta, Kazutomo Yoshii, Henry Hoffman, Allen Malony, Martin Schulz, and Pete Beckman. *Systemwide Power Management with Argo*. In HPPAC '16: Proceedings of the Twelfth Workshop on High-Performance, Power-Aware Computing, Chicago, IL, USA, May 2016.
- [5] Huiwei Lu, Sangmin Seo, and Pavan Balaji. *MPI+ULT: Overlapping Communication and Computation with User-Level Threads*. In HPCC '15: Proceedings of the 2015 IEEE 17th International Conference on High Performance Computing and Communications, pp. 444–454, New York, NY, USA, Aug. 2015, DOI:10.1109/HPCC-CSS-ICSS.2015.82.
- [6] Sangmin Seo, Robert Latham, Junchao Zhang, and Pavan Balaji. *Implementation and Evaluation of MPI Non-blocking Collective I/O*. In PPMM '15: Proceedings of the 2nd Workshop on Parallel Programming Model for the Masses, pp. 1084–1091, Shenzhen, Guangdong, China, May 2015, DOI:10.1109/CCGrid.2015.81.
- [7] Sangmin Seo, Jun Lee, Gangwon Jo, and Jaejin Lee. *Automatic OpenCL Work-Group Size Selection for Multi-core CPUs*. In PACT '13: Proceedings of the 22nd International Conference on Parallel Architectures and Compilation Techniques, pp. 387–397, Edinburgh, Scotland, UK, Sep. 2013, DOI:10.1109/PACT.2013.6618834. (36/208, 17.3%)
- [8] Jungwon Kim, Sangmin Seo, Jun Lee, Jeongho Nah, Gangwon Jo, and Jaejin Lee. *SnuCL: an OpenCL Framework for Heterogeneous CPU/GPU Clusters*. In ICS '12: Proceedings of the 26th International Conference on Supercomputing, pp. 341–352, San Servolo Island, Venice, Italy, Jun. 2012, DOI:10.1145/2304576.2304623. (36/161, 22.4%)
- [9] Choonki Jang, Jun Lee, Sangmin Seo, and Jaejin Lee. *An Automatic Code Overlaying Technique for Multicores with Explicitly-Managed Memory Hierarchies*. In CGO '12: Proceedings of the 2012 International Symposium on Code Generation and Optimization, pp. 219–229, San Jose, CA, USA, Mar. 2012, DOI:10.1145/2259016.2259045. (26/90, 28.9%)

- [10] Sangmin Seo, Gangwon Jo, and Jaejin Lee. *Performance Characterization of the NAS Parallel Benchmarks in OpenCL*. In IISWC '11: Proceedings of the 2011 IEEE International Symposium on Workload Characterization, pp. 137–148, Austin, TX, USA, Nov. 2011, DOI:[10.1109/IISWC.2011.6114174](https://doi.org/10.1109/IISWC.2011.6114174). (20/50, 40.0%)
- [11] Sangmin Seo, Junghyun Kim, and Jaejin Lee. *SFMalloc: A Lock-Free and Mostly Synchronization-Free Dynamic Memory Allocator for Manycores*. In PACT '11: Proceedings of the 20th International Conference on Parallel Architectures and Compilation Techniques, pp. 252–262, Galveston Island, TX, USA, Oct. 2011, DOI:[10.1109/PACT.2011.57](https://doi.org/10.1109/PACT.2011.57). (36/221, 16.3%)
- [12] Jun Lee, Jungwon Kim, Junghyun Kim, Sangmin Seo, and Jaejin Lee. *An OpenCL Framework for Homogeneous Manycores with no Hardware Cache Coherence*. In PACT '11: Proceedings of the 20th International Conference on Parallel Architectures and Compilation Techniques, pp. 56–67, Galveston Island, TX, USA, Oct. 2011, DOI:[10.1109/PACT.2011.12](https://doi.org/10.1109/PACT.2011.12). (36/221, 16.3%)
- [13] Jungwon Kim, Sangmin Seo, Jun Lee, Jeongho Nah, Gangwon Jo, and Jaejin Lee. *OpenCL as a Programming Model for GPU Clusters*. In LCPC '11: Proceedings of the 24th International Workshop on Languages and Compilers for Parallel Computing, pp. 76–90, Fort Collins, CO, USA, Sep. 2011, DOI:[10.1007/978-3-642-36036-7_6](https://doi.org/10.1007/978-3-642-36036-7_6).
- [14] Junghyun Kim, Sangmin Seo, and Jaejin Lee. *An Efficient Software Shared Virtual Memory for the Single-chip Cloud Computer*. In APSys '11: Proceedings of the 2nd ACM SIGOPS Asia-Pacific Workshop on Systems, pp. 17–21, Shanghai, China, Jul. 2011, DOI:[10.1145/2103799.2103804](https://doi.org/10.1145/2103799.2103804). (19/57, 33.3%)
- [15] Jaejin Lee, Jungwon Kim, Sangmin Seo, Seungkyun Kim, Jungho Park, Honggyu Kim, Thanh Tuan Dao, Yongjin Cho, Sung Jong Seo, Seung Hak Lee, Seung Mo Cho, Hyo Jung Song, Sang-Bum Suh, and Jong-Deok Choi. *An OpenCL Framework for Heterogeneous Multicores with Local Memory*. In PACT '10: Proceedings of the 19th International Conference on Parallel Architectures and Compilation Techniques, pp. 193–204, Vienna, Austria, Sep. 2010, DOI:[10.1145/1854273.1854301](https://doi.org/10.1145/1854273.1854301). (46/266, 17.3%)
- [16] Jaejin Lee, Jun Lee, Sangmin Seo, Jungwon Kim, Seungkyun Kim, and Zehra Sura. *COMIC++: A Software SVM System for Heterogeneous Multicore Accelerator Clusters*. In HPCA '10: Proceedings of the 16th IEEE International Symposium on High Performance Computer Architecture, pp. 329–340, Bangalore, India, Jan. 2010, DOI:[10.1109/HPCA.2010.5416633](https://doi.org/10.1109/HPCA.2010.5416633). (32/175, 18.3%)
- [17] Sangmin Seo, Jaejin Lee, and Zehra Sura. *Design and Implementation of Software-Managed Caches for Multicores with Local Memory*. In HPCA '09: Proceedings of the 15th IEEE International Symposium on High Performance Computer Architecture, pp. 55–66, Raleigh, NC, USA, Feb. 2009, DOI:[10.1109/HPCA.2009.4798237](https://doi.org/10.1109/HPCA.2009.4798237). (35/184, 19.0%)
- [18] Jaejin Lee, Sangmin Seo, Chihun Kim, Junghyun Kim, Posung Chun, Zehra Sura, Jungwon Kim, and SangYong Han. *COMIC: A Coherent Shared Memory Interface for Cell BE*. In PACT '08: Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques, pp. 303–314, Toronto, Canada, Oct. 2008, DOI:[10.1145/1454115.1454157](https://doi.org/10.1145/1454115.1454157). (30/159, 18.9%)

3.3. Technical Reports

- [1] Sangmin Seo, Abdelhalim Amer, Pavan Balaji, Cyril Bordage, George Bosilca, Alex Brooks, Adrián Castelló, Damien Genet, Thomas Herault, Prateek Jindal, Laxmikant V. Kalé, Sriram Krishnamoorthy, Jonathan Lifflander, Huiwei Lu, Esteban Meneses, Marc Snir, Yanhua Sun, and Pete Beckman. *Argobots: A Lightweight, Low-Level Threading and Tasking Framework*. ANL/MCS-P5515-0116, Argonne National Laboratory, 2016. [online] Available at: <http://www.mcs.anl.gov/papers/P5515-0116.pdf>

3.4. Poster Papers

- [1] Swann Perarnau, Rinku Gupta, Pete Beckman, Pavan Balaji, Cyril Bordage, George Bosilca, David Callahan, Franck Cappello, Jack Dongarra, Daniel Ellsworth, Brain Van Essen, Damien Genet, Roberto Gioiosa, Maya Gokhale, Thomas Herault, Henry Hoffman, Laxmikant Kale, Gokcen Kestor, Sriram Krishnamoorthy, Kamil Iskra, Edgar Leon, Jonathan Lifflander, Huiwei Lu, Allen Malony, Nikita Mishra, Kenneth Raffanetti, Barry Rountree, Martin Schulz, Sangmin Seo, Sameer Shende, Marc Snir, Wyatt Spear, Yanhua Sun, Rajeev Thakur,

- Kazutomo Yoshii, Xuechen Zheng, Huazhe Zhang, Judiceal Zounmevo. *ARGO: An Exascale Operating System and Runtime*. Poster presentation in SC '15: 2015 International Conference for High Performance Computing, Networking, Storage and Analysis, Austin, TX, US, Nov. 2015.
- [2] Jintao Meng, Yanjie Wei, Sangmin Seo, and Pavan Balaji. *SWAP-Assembler 2: Scalable Genome Assembler towards Millions of Cores - Practice and Experience*. Doctoral symposium in CCGrid '15: Proceedings of the 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, pp. 769–772, Shenzhen, Guangdong, China, May 2015, DOI:[10.1109/CCGrid.2015.128](https://doi.org/10.1109/CCGrid.2015.128).
 - [3] Wesley Bland, Huiwei Lu, Sangmin Seo, and Pavan Balaji. *Lessons Learned Implementing User-Level Failure Mitigation in MPICH*. Poster presentation in CCGrid '15: Proceedings of the 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, pp. 1123–1126, Shenzhen, Guangdong, China, May 2015, DOI:[10.1109/CCGrid.2015.51](https://doi.org/10.1109/CCGrid.2015.51).
 - [4] Swann Perarnau, Judicael Zounmevo, Sangmin Seo, Huiwei Lu, Kenneth Raffanetti, Rinku Gupta, Kamil Iskra, Kazutomo Yoshii, Pavan Balaji, Franck Cappello, Rajeev Thakur, Marc Snir, and Pete Beckman. *Argo: An Exascale Operating System and Runtime*. In the 4th Greater Chicago Area System Research Workshop (GCASR) 2015, Chicago, IL, USA, Apr. 2015.
 - [5] Jungwon Kim, Sangmin Seo, Jun Lee, Jeongho Nah, Gangwon Jo, and Jaejin Lee. *OpenCL as a Unified Programming Model for Heterogeneous CPU/GPU Clusters*. Poster presentation in PPOPP '12: Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pp. 299–300, New Orleans, LA, USA, Feb. 2012, DOI:[10.1145/2145816.2145863](https://doi.org/10.1145/2145816.2145863).
 - [6] Jun Lee, Sangmin Seo, and Jaejin Lee. *A Software-SVM-based Transactional Memory for Multicore Accelerator Architectures with Local Memory*. Poster presentation in PACT '10: Proceedings of the 19th International Conference on Parallel Architectures and Compilation Techniques, pp. 567–568, Vienna, Austria, Sep. 2010, DOI:[10.1145/1854273.1854355](https://doi.org/10.1145/1854273.1854355).

3.5. Publications in Korean

- [1] Junghyun Kim, Jungho Park, Gangwon Jo, Thanh Tuan Dao, Jingyoung Joo, Jaehoon Jung, Jungwon Kim, Sangmin Seo, Jun Lee, Jeongho Nah, and Jaejin Lee. *SnuCL : OpenCL Programming Environment for Heterogeneous Manycore Clusters*. Communications of KIISE, Vol. 32, No. 5, pp. 66–76, May 2014.
- [2] Gangwon Jo, Sangmin Seo, Jeongho Nah, Jungwon Kim, Junghyun Kim, Jun Lee, Jungho Park, Yong-Jun Lee, Hongjune Kim, Sooyeon Kang, Jinyoung Joo, Seonmyeong Park, Wookeun Jung, Kihyun Im, and Jaejin Lee. *Trends on Heterogeneous Supercomputers and a Case Study on the Development of a Supercomputer Chundoong*. Communications of KIISE, Vol. 31, No. 4, pp. 34–41, Apr. 2013.
- [3] Junghyun Kim, Sangmin Seo, Jaejin Lee, Mikyoung Park, Jongyoung Lee, and Taekyeong Ko. *The Design and Implementation of a Cache Simulator for Multicore Systems*. Korea Computer Congress, Jun. 2009.
- [4] Taejun Ha, Sangmin Seo, Posung Chun, and Jaejin Lee. *Automatic Detection of Memory Subsystem Parameters for Embedded Systems*. Journal of KIISE: Computing Practices and Letters, Vol. 20, No. 5, pp. 350–354, May 2009.
- [5] Taejun Ha, Sangmin Seo, Posung Chun, and Jaejin Lee. *Automatic Detection of Memory Subsystem Parameters for Embedded Systems*. Korea Computer Congress, Jun. 2008.

3.6. Presentation at Major Conferences and Symposia

- [1] *Argobots and its Application to Charm++*. The 14th Annual Workshop on Charm++ and its Applications, Champaign, IL, USA, Apr. 2016.
- [2] *Implementation and Evaluation of MPI Nonblocking Collective I/O*. The 2nd Workshop on Parallel Programming Model for the Masses (PPMM), Shenzhen, Guangdong, China, May 2015.
- [3] *Automatic OpenCL Work-Group Size Selection for Multicore CPUs*. The 22nd International Conference on Parallel Architectures and Compilation Techniques (PACT), Edinburgh, Scotland, UK, Sep. 2013.

- [4] *Performance Characterization of the NAS Parallel Benchmarks in OpenCL*. The 2011 IEEE International Symposium on Workload Characterization (IISWC), Austin, TX, USA, Nov. 2011,
- [5] *SFMalloc: A Lock-Free and Mostly Synchronization-Free Dynamic Memory Allocator for Manycores*. The 20th International Conference on Parallel Architectures and Compilation Techniques (PACT), Galveston Island, TX, USA, Oct. 2011.
- [6] *Design and Implementation of Software-Managed Caches for Multicores with Local Memory*. The 15th IEEE International Symposium on High Performance Computer Architecture (HPCA), Raleigh, NC, USA, Feb. 2009.

3.7. Invited Talks

- [1] *BOLT: OpenMP over Lightweight Threads*. OpenMP Booth (#611) Talk @ SC 2016, Salt Lake City, UT, USA, Nov. 2016.
- [2] *User-Level Threads and OpenMP*. 4th XcalableMP Workshop, Tokyo, Japan, Nov. 2016.
- [3] *BOLT: OpenMP over Lightweight Threads*. The 5th Joint Laboratory for Extreme-Scale Computing (JLESC) Workshop, Lyon, France, Jun. 2016.
- [4] *Enhancing OpenMP and MPI with Lightweight Threads*. Seoul National University, Seoul, Korea, Mar. 2016.
- [5] *Enhancing OpenMP and MPI with Lightweight Threads*. Sungkyunkwan University, Suwon, Korea, Mar. 2016.
- [6] *MPI and OpenMP with User-Level Threads*. Intel HPC Developer Conference, Beijing, China, Jan. 2016.
- [7] *MPI and OpenMP with User-Level Threads*. National Supercomputer Center in Jinan, Jinan, China, Jan. 2016.
- [8] *Performance Analysis and Optimizations of Argobots*. The 4th Joint Laboratory for Extreme-Scale Computing (JLESC) Workshop, Bonn, Germany, Dec. 2015.
- [9] *Argobots and Its User-defined Schedulers*. COOLR Annual Meeting, Argonne National Laboratory, Lemont, IL, USA, Sep. 2015.
- [10] *Argobots: Lightweight Low-level Threading/Tasking Framework*. The 3rd Joint Laboratory for Extreme-Scale Computing (JLESC) Workshop, Barcelona, Spain, Jun. 2015.
- [11] *Argobots: Lightweight Low-level Threading/Tasking Framework*. Tongji-ANL Workshop, Tongji University, China, May 2015.
- [12] *An OpenCL Framework for Multicore CPUs*. Argonne National Laboratory, Lemont, IL, USA, Feb. 2014.
- [13] *Intel Xeon Phi Coprocessor Programming*. Hankook Tire Co., LTD., Daejeon, Korea, Oct. 2013.
- [14] *Intel Xeon Phi Coprocessor Programming*. Samsung SDI, Cheonan, Korea, Aug. 2013.
- [15] *Introduction to OpenCL Programming*. National Supercomputing Summer School @ UNIST, Ulsan, Korea, Jul. 2013.
- [16] *Introduction to OpenCL Programming*. Accelerator Programming Winter School, Siheung, Korea, Feb. 2013.
- [17] *HPC Technology Trend in Korea*. HP-CAST Korea 2013, Seoul, Korea, Apr. 2013.

3.8. Software Contributed

- [1] **BOLT**: OpenMP over Lightweight Threads
- <http://www.bolt-omp.org/>
- [2] **Argobots**: A Lightweight Low-level Threading and Tasking Framework
- <http://www.argobots.org/>
- [3] **MPICH**: A High-Performance, Portable Implementation of MPI
- <http://www.mpich.org/>
- [4] **SnuCL**: An OpenCL Framework for Heterogeneous Clusters
- <http://snuc1.snu.ac.kr/>
- [5] **SNU NPB Suite**

- <http://aces.snu.ac.kr/software/snu-npb/>
- [6] **SNU-SAMSUNG OpenCL Framework**
- <http://aces.snu.ac.kr/software/snu-samsung-opencl/>
- [7] **SFMalloc**
- <http://aces.snu.ac.kr/software/sfmalloc/>
- [8] **COMIC: Coherent Shared Memory Interface for Cell BE**
- <http://aces.snu.ac.kr/software/comic/>

3.9. Patents

- [1] Jaejin Lee and Sangmin Seo. *Method and System for Determining Work-Group Size and Computer Readable Recording Medium Therefor*. Korean Patent 10-1537725 (Application No. 10-2013-0166372), Filed December 30, 2013, Issued Jul. 13, 2015.

3.10. Other

3.10.1. International Standards

- [1] *MPI: A Message-Passing Interface Standard, Version 3.1*. The Message Passing Interface Forum, Jun. 4th, 2015.

3.10.2. Demos

- [1] *ARGO: An Exascale Operating System and Runtime*. DOE Booth (#502) @ SC 2015, Nov. 17, 2015.

4. Research Grants

- [1] **PI: BOLT: OpenMP over Lightweight Threads**. Argonne National Laboratory, Laboratory Directed Research and Development (LDRD) Innovate program. Period: October 2016 – September 2018. Amount (single institute grant): \$420,000.
- [2] **Co-PI** (PI: Jaejin Lee, Seoul National University): *Development of the Big Data Processing Platform based on Manycore Performance Acceleration*. Small and Medium Business Administration, Period: July 2013 – June 2014. Amount: 89,827,000 KRW.
- [3] **PI: Chundoong-S: High-performance Accelerator Supercomputing Solution**. Small and Medium Business Administration, Period: July 2013 – February 2014. Amount: 35,000,000 KRW.

5. Professional Activities

5.1. Conference/Workshop Organizing Committee

- [1] Program Co-chair (with Antonio J. Peña, Barcelona Supercomputing Center, Spain). *International Workshop on Accelerators and Hybrid Exascale Systems (AsHES)*, 2017.
- [2] Finance and Registrations Chair. *European/USA MPI Users' Group Meeting (EuroMPI/USA)*, 2017.
- [3] Posters Chair. *International Conference on Parallel Processing (ICPP)*, 2017.
- [4] Web/Publicity Chair. *International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, 2017.
- [5] Finance and Registration Co-chair (with Jose Daniel Garcia, Universidad Carlos III de Madrid, Spain). *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, 2017.
- [6] Finance and Registration Co-chair (with Mary Dzieliski, Argonne National Laboratory, USA). *IEEE International Conference on Cluster Computing (Cluster)*, 2015.
- [7] Web Co-chair (with Wenqi Sun, Tsinghua University, China). *International Conference on Parallel Processing (ICPP)*, 2015.
- [8] Finance and Registration Co-chair (with Mary Dzieliski, Argonne National Laboratory, USA). *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, 2015.

5.2. Conference/Workshop Program Committee

- [1] International Supercomputing Conference (ISC) 2017 - Programming Models and Systems Software track
- [2] IEEE International Parallel & Distributed Processing Symposium (IPDPS) 2017 - Software track
- [3] European/USA MPI Users' Group Meeting (EuroMPI/USA) 2017
- [4] International Conference on Cloud Computing and Big Data (CCBD): 2015, 2016
- [5] International Workshop on Fault Tolerant Systems (FTS): 2015, 2016
- [6] International Workshop on Scheduling and Resource Management for Parallel and Distributed Systems (SRM-PDS): 2015, 2016
- [7] International Workshop on Accelerators and Hybrid Exascale Systems (AsHES): 2016

5.3. Working Groups

OpenMP Language Committee	SEP 2016 - present
MPI Forum	JUN 2014 - present

5.4. Teaching

<i>Guest Lecturer</i>	SEP 21, 2016
<ul style="list-style-type: none">- CS546 Parallel and Distributed Processing, Illinois Institute of Technology, Chicago, IL, USA- Lecture: Introduction to OpenMP	
<i>Lecturer</i>	JUN 2013 - NOV 2013
<ul style="list-style-type: none">- SMBA R&D Academy, Seoul National University, Seoul, Korea- Course: Parallel Programming Expert Course using the Supercomputer 'Chundoong'- Responsible for teaching parallel programming concepts and skills with Pthreads, OpenMP, OpenCL, and MPI	
<i>Teaching Assistant</i>	
<ul style="list-style-type: none">- 4190.409 Compilers, Spring 2008, Seoul National University, Seoul, Korea- 4190.409 Compilers, Spring 2011, Seoul National University, Seoul, Korea- Responsible for providing advice upon students' questions and grading homework and exams	
<i>Lecturer</i>	JUN 2008
<ul style="list-style-type: none">- Bit Academy, Seoul, Korea- Responsible for two-week course where students understand the concept of Component Object Model (COM) and develop programming skill using COM and ActiveX	

5.5. Students Mentored

Hoang-Vu Dang	MAY 2016 - AUG 2016
<ul style="list-style-type: none">- Ph.D. candidate, University of Illinois at Urbana-Champaign, USA.- Topic: Communication-Aware Thread Scheduling in MPI.	
Adrián Castelló	MAY 2015 - AUG 2015
<ul style="list-style-type: none">- Ph.D. student, Universitat Jaume I, Spain.- Topic: Improving the Performance of OpenMP using Lightweight Threads.- Currently serving as his Ph.D. thesis committee	
Jintao Meng	SEP 2014 - OCT 2014
<ul style="list-style-type: none">- Ph.D. candidate, Shenzhen Institutes of Advanced Technology, China.- Topic: SWAP-Assembler: Scalable Genome Assembler towards Millions of Cores Practice and Experience.	
Priyanka Ghosh	MAY 2014 - AUG 2014
<ul style="list-style-type: none">- Ph.D. student, University of Houston, USA.- Topic: Minimizing Cost of Data Movement by Optimizing Data-reuse on Large Scale GA-based Applications.	

5.6. Others

5.6.1. Session Chair

- [1] Birds-of-a-Feather @ SC 2015: Towards Standardized, Portable and Lightweight User-Level Threads and Tasks (co-chaired with Pavan Balaji, Argonne National Laboratory, USA)
- [2] Session 9B: Distributed Resource Management @ CCGrid 2015

5.6.2. Journal Reviewer

- [1] IEEE Transactions on Parallel and Distributed Systems (TPDS): 2016
- [2] IEEE Transactions of Cloud Computing (TCC): 2014–2016
- [3] Elsevier Parallel Computing (ParCo): 2014, 2016
- [4] Journal of Concurrency and Computation: Practice and Experience (CCPE): 2014–2016
- [5] International Journal of High Performance Computing (IJHPCA): 2016
- [6] Journal of Parallel and Distributed Computing (JPDC): 2016

5.6.3. Conference/Workshop External Reviewer

- [1] International Workshop on Extreme Scale Programming Models and Middleware (ESPM2): 2015, 2016
- [2] IFIP International Conference on Network and Parallel Computing (NPC): 2016
- [3] Hawaii International Conference on System Sciences (HICSS): 2016, 2017
- [4] IEEE Annual Symposium on High-Performance Interconnects (HotI): 2014, 2016
- [5] European MPI Users' Group Meeting (EuroMPI): 2015, 2016
- [6] ACM International Conference on Supercomputing (ICS): 2012, 2013, 2016
- [7] International Workshop on Programming Models and Applications for Multicores and Manycores (PMAM): 2016
- [8] IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid): 2015, 2016
- [9] IEEE International Parallel & Distributed Processing Symposium (IPDPS): 2016
- [10] International Workshop in Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS): 2015
- [11] IEEE International Conference on Parallel and Distributed Systems (ICPADS): 2014, 2015
- [12] IEEE International Conference on High Performance Computing (HiPC): 2014, 2015
- [13] International Conference on Parallel Architectures and Compilation Techniques (PACT): 2011, 2012, 2015
- [14] International European Conference on Parallel and Distributed Computing (Euro-Par): 2015
- [15] IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS): 2013